

CONTENT DISPLAY MONITOR

Trevor I. Blumenau

BACKGROUND OF THE INVENTION

1. Field of the Invention

5 ~~This invention relates to monitoring the display and~~  
~~observation of content by a computer system.~~ The invention  
also relates to monitoring the display and observation at a  
content display site of content that is provided by a content  
provider site over a network to the content display site. The  
10 invention further relates to the provision of updated and/or  
tailored content from a content provider site to a content  
display site so that the content provider's current content is  
always displayed at the content display site.

2. Related Art

15 A large amount of human activity consists of the  
dissemination of information by information providers (content  
providers) to information consumers (observers). Recently,  
computer networks have become a very popular mechanism for  
accomplishing information dissemination. The use of computer  
20 networks for information dissemination has necessitated or  
enabled new techniques to accomplish particular functions  
related to the dissemination of information.

For example, information providers of all types have an  
interest in knowing the extent and nature of observation of the  
25 information that they disseminate. Information providers that  
disseminate information over computer networks also have this  
interest. However, the use of networked computers for  
information dissemination can make it difficult to ascertain  
who is observing the disseminated information and how, since  
30 information can be accessed rapidly from a remote location by  
any of a large number of possible observers whose identity is

often not predictable beforehand, and since control over the display of the information once disseminated may not be possible, practical or desirable.

Among information providers, advertisers have particular  
5 interest in knowing how and to what extent their advertisements  
are displayed and/or observed, since such knowledge can be a  
key element in evaluating the effectiveness of their  
advertising and can also be the basis for payment for  
advertising. Mechanisms for obtaining such information have  
10 been developed for advertisements disseminated in conventional  
media, e.g., audiovisual media such as television and radio,  
and print media such as magazines and newspapers. For example,  
the well-known Nielsen television ratings enable advertisers to  
15 gauge the number of people that likely watched advertisements  
during a particular television program. As advertising over a  
computer network becomes more common, the importance of  
developing mechanisms for enabling advertisers to monitor the  
display and observation of their advertisements disseminated  
over a computer network increases.

20 Previous efforts to monitor the display of advertising (or  
other content) disseminated over a computer network have been  
inadequate for a variety of reasons, including the limited  
scope of the monitoring information obtained, the ambiguous  
nature of the monitoring information, the incompleteness of the  
25 monitoring information, and the susceptibility of the  
monitoring information to manipulation. Review of some of the  
techniques that have previously been used to acquire monitoring  
information regarding the display of content (e.g.,  
advertising) disseminated over a particular computer network -  
30 the World Wide Web portion of the Internet computer network -  
will illustrate the deficiencies of existing techniques for  
monitoring the display of content disseminated over a computer  
network.

FIGS. 1A and 1B are simplified diagrams of a network illustrating operation of a previous system for monitoring requests for content over the World Wide Web. In FIGS. 1A and 1B, a content provider site 101 (which can be embodied by, for example, a server computer) can communicate with a content display site 102 (which can be embodied by, for example, a client computer) over the network communication line 103. The server computer at the content provider site 101 can store content colloquially referred to as a "Web page." The client computer at the content display site 102 executes a software program, called a browser, that enables selection and display of a variety of Web pages stored at different content provider sites. When an observer at the content display site 102 wishes to view a particular Web page, the observer causes the client computer at the content display site 102 to send a request to the appropriate server computer, e.g., the server computer at the content provider site 101, as shown in FIG. 1A. The server computers at content provider sites all include a software program (in the current implementation of the World Wide Web, this is an http daemon) that watches for such incoming communications. Upon receipt of the request, the server computer at the content provider site 101 transfers a file representing the Web page (which, in the current implementation of the World Wide Web, is an html file) to the client computer at the content display site 102, as shown in FIG. 1B. This file can itself reference other files (that may be stored on the server computer at the content provider site 101 and/or on other server computers) that are also transferred to the content display site 102. The browser can use the transferred files to generate a display of the Web page on the client computer at the content display site 102. The http daemon, in addition to initiating the transfer of the appropriate file or files to the content display site 102, also makes a record of

requests for files from the server computer on which the daemon resides. The record of such requests is stored on the server computer at the content provider site 101 in a file 104 that is often referred to as a "log file."

5       The exact structure and content of log files can vary somewhat from server computer to server computer. However, generally, log files include a list of transactions that each represent a single file request. Each transaction includes multiple fields, each of which are used to store a predefined  
10 type of information about the file request. One of the fields can be used to store an identification of the file requested. Additional fields can be used to store the IP (Internet Protocol) address of the client computer that requested the particular file, the type of browser that requested the file,  
15 a time stamp for the request (i.e., the date and time that the request was received by the server computer), the amount of time required to transfer the requested file to the client computer, and the size of the file transferred. Other information about file requests can also be stored in a log  
20 file.

Previous methods for monitoring the display of content distributed over the World Wide Web have used the information stored in the log file. For example, one previous method has consisted of simply determining the number of transactions in  
25 the log file and counting each as a "hit" on a Web page, i.e., a request for a Web page. The number of hits is deemed to approximate the number of times that the Web page has been viewed and, therefore, the degree of exposure of the content of the Web page to information consumers.

30       There are a number of problems with this approach however. For example, as indicated above, a request for a Web page may cause, in addition to the request for an initial html file, requests for other files that are necessary to generate the Web

page. If these other files reside on the same server computer as the initial html file, additional transactions are recorded in the log file. Thus, a request for a single Web page can cause multiple transactions to be recorded in the log file. As 5 can be appreciated, then, the number of times that a Web page is transferred to a content display site can be far less than the number of transactions recorded in the log file. Moreover, without further analysis, there is no way to accurately predict the relationship between the number of transactions in the log 10 file and the number of times that a Web page has been transferred to the content display site. Such inaccuracy can be very important to, for example, advertisers - whose cost of advertising is often proportional to the measured exposure of the advertising - since the measured exposure of their 15 advertising (and, thus, its cost) may be based upon the number of hits on a Web page containing their advertisement.

A method to overcome this problem has been used. By analyzing the contents of the log file to determine which file was requested in each transaction, it may be possible to 20 differentiate transactions in which the initial html file needed to generate a Web page is requested from transactions in which the requested file is one which is itself requested by another file, thus enabling "redundant" transactions to be identified and eliminated from the hit count. While such an 25 approach can increase the accuracy of counting Web page hits, it still suffers from several problems.

For example, log file analysis may result in some undercounting of Web page hits, apart from any overcounting. This is because, once transferred to a client computer at a 30 content display site, the files necessary to generate a Web page can be stored ("cached") on that client computer, thus enabling an observer at the content display site to view the Web page again without causing the client computer to make

another request to the content provider server computer from which the Web page was initially retrieved. Consequently, the observer can view the Web page without causing transactions to be added to the log file, resulting in undercounting of the number of Web page hits.

Additionally, log files are subject to manipulation, either directly or indirectly. For example, an unscrupulous content provider could directly manipulate the log file by retrieving and editing the log file to add phony transactions, thus artificially increasing the number of Web page hits and making the Web page appear to be more popular than it really is. This problem can be ameliorated by causing the log files to be transferred periodically at predetermined times (e.g., each night at 12:00 midnight) from the server computer at the content provider site to a neutral network site; however, the log file can still be manipulated during the time between transfers.

A log file might be manipulated indirectly, for example, by programming one or more computers to continually request a Web page, thereby generating a large number of hits on that Web page. While the log file would contain transactions corresponding to actual file requests associated with the Web page, these requests would be artificial requests that would almost certainly not result in a display of the Web page, and certainly not in the observation of the Web page. Moreover, checking the contents of the log file for an unusually high number of requests from a particular IP address (i.e., client computer) may not enable such manipulation to be detected, since a large number of requests may legitimately come from a client computer that serves many users (for example, the proprietary network America Online™ has a handful of computers that are used by many users of that network to make connection to the Internet and World Wide Web).

It may be possible to identify the real origin of requests for content using "cookies." A cookie enables assignment of a unique identifier to each computer from which requests really emanate by transferring the identifier to that computer with content transferred to that computer. Future requests for content carry this identifier with them. The identifier can be used, in particular, to aid in identification of indirect log file manipulation, as described above, and, more generally, to enable more robust log file analysis.

10 Notwithstanding such enhancement, cookies do not overcome a fundamental problem with the use and analysis of log files to ascertain information regarding the display of content provided over the World Wide Web. That is, as highlighted by the overcounting problem associated with the above-described  
15 artifice and the undercounting problem associated with caching of content at the content display site, log files only store information about file requests. A log file does not even indicate whether the requested file was actually transferred to the requesting client computer (though, typically, such file  
20 transfer would occur). Nor does a log file include any information about how the file was used once transferred to the requesting client computer. In particular, log files do not provide any information regarding whether the content represented by the requested file is actually displayed by the  
25 client computer at the content display site, much less information from which conclusions can be deduced regarding whether - and if so, how - the content was observed by an observer. These limitations associated with the content of a log file cannot be overcome by a monitoring approach based on  
30 log file analysis. Moreover, log file analysis is calculation intensive, requiring hours in some instances to extract the desired information from the log file.

Another method of monitoring the display of content

disseminated over the World Wide Web uses an approach similar to that of the Nielsen ratings system used in monitoring television viewing. In this method, monitoring software is added to the browser implemented on the client computers of a selected number of defined observers (e.g., families) to enable acquisition of data regarding advertising exposure on those computers. This information is then used to project patterns over the general population.

However, this approach also has several disadvantages.

10 First, only a limited amount of data is collected, i.e., data is only obtained regarding a small number of information consumers. As with any polling method, there is no guarantee that the data acquired can be extrapolated to the general population, even if the observers selected for monitoring are

15 chosen carefully and according to accepted sampling practices. Second, as the size of the World Wide Web (or other computer network for which this method is used) grows, i.e., as the number of content provider sites increases, the number of monitored observers necessary to ensure accurate representation

20 of the usage of all content provider sites must increase, since otherwise there may be few or no observer interactions with some content provider sites upon which to base projections. It may not be possible to find an adequate number of appropriate observers to participate in the monitoring process,

25 particularly given concerns with the attendant intrusion into the privacy of the selected observers. Third, installation of the monitoring software on a client computer to be compatible with a browser presents a number of problems. Such installation requires active participation by observers; since

30 observers typically do not reap benefit from operation of the monitoring software, they may be reluctant to expend the effort to effect installation. The monitoring software must continually be revised to be compatible with new browsers and

new versions of old browsers. To enable monitoring of a large number of client computers, the software must be tested for compatibility with a wide variety of computing environments. And, as currently implemented, such monitoring software is also  
5 dependent upon the computing platform used, making it necessary to revise the monitoring software for use with new computing platforms or risk skewing the demographics of the sample users.

In addition to desiring information regarding the display and observation of the content that they provide, content  
10 providers also often desire to provide content to a content display site that is particularly tailored for observation (e.g., according to various demographic characteristics of an expected observer) at that content display site. For example, text content should be expressed in a language that the  
15 observer can understand. If appropriate for the content, it is desirable to tailor the content according to, for example, the age, sex or occupation of the observer.

Such tailoring of content has previously been enabled by modifying the http daemon on a computer at the content provider  
20 site to cause a particular version of a set of content to be transferred to a requesting content display site based upon the IP address of that content display site. While such tailoring of content is useful, it is desirable to be able to tailor the presentation of content in additional ways not  
25 enabled by this approach.

Content providers also often desire to provide their content with the content of other content providers. For example, it is a common practice for content providers (referred to here as "primary content providers") on the World  
30 Wide Web to include advertisements from other entities (referred to here as "secondary content providers") as part of the content provider's Web page. In such situations, it is desirable for the secondary content provider to be able to

03240364460  
easily update and/or appropriately tailor (e.g., according to characteristics of the requester) the content that they supply to the primary content provider. This could be accomplished by causing the primary content provider site to contact the  
5 secondary content provider site - each time that the primary content provider receives a request for content that includes the secondary content - to retrieve the secondary content (thus ensuring that updated, appropriately tailored secondary content is used) or check whether updated or tailored secondary content  
10 is available (if so, the content is retrieved). (This method could also be modified so that content retrieval or a check for updated and/or tailored content is only performed according to a predetermined schedule.) However, both the primary content provider and the secondary content provider may not want their  
15 systems burdened with the extra computational capacity required to handle the multitude of requests that would be needed to effect this operation. Alternatively, the primary content provider could collect and store the updated and tailored content from the secondary content providers at the primary  
20 content provider site. However, the burden associated with collecting and managing the content from secondary content providers may be more than the primary content provider wants to shoulder.

One way that this functionality can be achieved without  
25 creating an undesirable burden on the primary or secondary content providing systems is by providing a secondary content storage site that can continually store the most recent content provided by a secondary content provider, as well as different sets of content tailored for particular situations (e.g.,  
30 display by particular observers or at particular times). FIGS. 2A through 2D are simplified diagrams of a network illustrating the operation of such a system. In FIG. 2A, a content display site 202 makes a request over the network

communication line 203 to the primary content provider site 201 for content that includes the secondary content. In FIG. 2B, the primary content provider site 201 transfers the file or files stored at the primary content provider site 201 that are necessary to generate a display of the primary content. These files include appropriate reference to a file or files stored at a secondary content storage site 204 that includes the most updated and/or appropriately tailored secondary content for display with the primary content. As shown in FIG. 2C, this reference causes the content display site 202 to request the secondary content from the secondary content storage site 204. In FIG. 2D, the secondary content is transferred from the secondary content storage site 204 to the content display site 202 for display at the content display site 202.

However, while this system can relieve the primary content provider of the burden of managing the acquisition, storage and provision of secondary content (a burden that can become rather onerous when many secondary content providers are providing content to the primary content provider), the system has a characteristic that can make it undesirable for many content providers. The secondary content storage site not only manages the secondary content, it also provides the secondary content when requests for primary content are made to the primary content provider. Moreover, the secondary content is frequently content, such as graphics files used to generate visual images (which frequently dominate advertisements), that has a high bandwidth requirement for transmission over the network. By taking control of the transmission of secondary content to the content display site, the secondary content storage site is also frequently taking control of the most bandwidth sensitive parts of the content provided by the primary content providers. The operator of the secondary content storage site may not provide a system that addresses

the bandwidth requirements to the satisfaction of the primary content provider, so that the presentation of the combined primary and secondary content occurs more slowly than desired by the primary content provider. Thus, this approach causes  
5 the primary content provider to lose control of a critical aspect of their operation.

#### SUMMARY OF THE INVENTION

The invention can enable monitoring of the display of content by a computer system. Moreover, the invention can  
10 enable monitoring of the content display to produce monitoring information from which conclusions may be deduced regarding the observation of the content display by an observer. The invention can also enable monitoring of the display at a content display site of content that is provided by a content  
15 provider site over a network to the content display site. Additionally, the invention can enable the expeditious provision of updated and/or tailored content over a network from a content provider site to a content display site so that the content provider's current and appropriately tailored  
20 content is always displayed at the content display site.

Aspects of the invention related to transfer of content over a network are generally applicable to any type of network. However, it is contemplated that the invention can be particularly useful with a computer network, including private  
25 computer networks (e.g., America Online™) and public computer networks (e.g., the Internet). In particular, the invention can be advantageously used with computer networks or portions of computer networks over which video and/or audio content are transferred from one network site to another network site for  
30 observation, such as the World Wide Web portion of the Internet. Additionally, the invention is particularly useful in monitoring the display of content obtained over such a

network using an interactive browser to acquire and view the content in real time.

In one aspect of the invention, the display of content by a computer system can be monitored by monitoring the position of the content display on a display screen of the computer system and evaluating the position of the content display on the display screen to produce monitoring information regarding display of the content. Monitoring of content display according to this aspect of the invention can be further enabled by monitoring the position of one or more other images on the display screen and comparing the position of the content display to the position of the other images to produce the monitoring information. In particular, this aspect of the invention can enable a determination as to whether (and for how long) the content display is hidden by one of the other images, and, further, whether the content display is fully hidden or partially hidden (and for how long the content display is fully and partially hidden, respectively). This information can be useful to, for example, indicate the amount of time that the content display was visible to an observer for observation, or to aid the content provider in determining in which regions of a display screen his content is most likely to be unobstructed. This aspect of the invention can also enable determination of the number of times that an on-screen pointer (e.g., a mouse arrow or a cursor) entered an area defined by the content display. This information may be useful in determining how attentive the observer was to the content, since an observer frequently watches the position of the on-screen pointer when viewing the display screen.

In another aspect of the invention, the display of content by a computer system can be monitored by monitoring the change in time of a characteristic of the content display and evaluating the change in time of the characteristic of the

content display to produce monitoring information regarding display of the content. Monitoring of content display according to this aspect of the invention can be further enabled by monitoring the change in time of a characteristic of the computer system and comparing the change in time of the characteristic of the content display to the change in time of the characteristic of the computer system to produce the monitoring information. This aspect of the invention can also enable, as discussed above, determination as to whether (and for how long) the content display is fully or partially hidden by another displayed image, as well as the number of times that an on-screen pointer entered an area defined by the content display.

In still another aspect of the invention, in a computer system in which the content is displayed in response to an instruction that is provided from a source external to the computer system and the system for monitoring (e.g., an instruction provided by a user of the computer system), the beginning and end of a display of the content can be ascertained so that monitoring of the display of content by the computer system can begin at the beginning of the content display and end at the end of the content display. The monitoring can occur in accordance with other aspects of the invention described herein. The monitoring can also determine the duration of the display of the content. Since the occurrence of monitoring according to this aspect of the invention is coincident with the display of the content to be monitored, the monitoring expends processing capability of the computer system only when necessary, while simultaneously assuring that monitoring occurs at all times that the content is displayed.

In yet another aspect of the invention, where content is provided by a content provider site over a network to a content

display site for display at the content display site, a mechanism for monitoring the display of the content can be transferred from the content provider site to the content display site in response to (e.g., together with) the transfer of content from the content provider site. The monitoring can occur in accordance with other aspects of the invention described herein. Monitoring information obtained regarding the display of the content at the content display site can be transferred to a remote site that is part of the network. The remote site can, but need not necessarily be, the content provider site from which the content was transferred to the content display site. Where the remote site is such content provider site, the monitoring information can then, in turn, be transferred from the content provider site to a second remote site. Further, where the remote site is such content provider site, the monitoring information can be transferred from the content display site to the content provider site via a communication means that is different from the communication means used to transfer the content from the content provider site to the content display site, a feature that can be useful, for example, when the network is the World Wide Web. This aspect of the invention provides a heretofore unavailable monitoring capability for obtaining information about how content is displayed on a network such as the World Wide Web. In particular, it has not previously been possible to monitor content transferred from a content provider site on the World Wide Web once the content has been transferred to a content display site.

In a still further aspect of the invention, in a network which operates according to a protocol that enables new content to be transferred to a content display site in response to selection of a portion of the content currently being displayed at the content display site, a mechanism for monitoring the

display of content can be transferred from a content provider site to a content display site so that the mechanism for monitoring operates at the content provider site. The monitoring can occur in accordance with other aspects of the invention described herein. Monitoring information obtained regarding the display of the content at the content display site can be transferred to a remote site that is part of the network. This aspect of the invention is particularly advantageous when at least some of the content being monitored comprises a graphical display. As discussed above with respect to the immediately preceding aspect of the invention, this aspect of the invention provides a heretofore unavailable monitoring capability for obtaining information about how content is displayed when retrieved over a network in an interactive browsing environment, such as occurs on the World Wide Web.

*insd2* *d2* ~~In yet a further aspect of the invention, the display of content that is provided by a content provider site over a network to the content display site is monitored to produce monitoring information, then the monitoring information is transferred to a remote site of the computer network that is different from the content provider site. According to this aspect, the monitoring information can first be transferred to the content providing site before eventual transfer to the remote site, so long as the monitoring information cannot be stored at the content provider site, or accessed or manipulated at the content provider site before transfer to the remote site. Access to the monitoring information at the remote site can be allowed to enable interaction with, but not modification of, the monitoring information. This aspect of the invention provides a system configuration that can overcome the problem of possible tampering with the substance of the monitoring information by the content provider. Further, this aspect of~~

the invention can be implemented so that the content and the monitoring instructions are stored at the remote site and transferred to the content display site when requested by an observer at the content display site, thus relieving the content provider of storing and managing the content and monitoring instructions at the content provider site and thereby simplifying use of the invention for the content provider.

In another aspect of the invention, current and/or tailored content can be provided to a content display site from a content provider site. The content can include both primary content (from the content provider site) and secondary content (provided by third parties). The primary and secondary content can be provided from a secondary content provider site to an application manager site. When the application manager site receives new content (in particular, updated and/or tailored content) from any content provider site, that content is transferred to content provider sites that use that content. Updated and/or tailored content is therefore available for transfer to a content display site immediately upon receipt of a request for the content from the content display site. This aspect of the invention relieves the primary content provider of the need to manage the storage of content, while reserving control over the provision of that content to the primary content provider, thereby enabling the content provider to ensure that the bandwidth requirement of the content provided from the content provider site are met.

#### BRIEF DESCRIPTION OF THE DRAWINGS

FIGS. 1A and 1B are simplified diagrams of a network illustrating operation of a previous system for monitoring requests for content over the World Wide Web.

FIGS. 2A, 2B, 2C and 2D are simplified diagrams of a network illustrating operation of a previous system for enabling retrieval of updated and/or tailored secondary content for use in primary content provided over the network.

5 FIGS. 3A, 3B and 3C are simplified diagrams of a network illustrating operation of one embodiment of the invention.

FIGS. 4A, 4B and 4C are simplified views of a display screen including a content display and other images, illustrating an unobstructed, fully hidden, and partially  
10 hidden content display, respectively.

FIG. 4D is a simplified view of a display screen including a content display and other images, illustrating a content display that is only partially hidden, but that would be determined to be fully hidden according to a method of the  
15 invention. FIG. 4E is a simplified view of the display screen shown in FIG. 4D, illustrating how another method of the invention can correctly determine the content display to be partially hidden.

FIG. 4F is a simplified view of a display screen including  
20 a content display and other images, illustrating a display that is partially hidden, but that may be determined to be unobstructed according to a method of the invention.

FIGS. 5A, 5B and 5C are simplified diagrams of a network illustrating operation of another embodiment of the invention.

25 FIGS. 6A, 6B, 6C and 6D are simplified diagrams of a network illustrating operation of still another embodiment of the invention.

*Insert*

DETAILED DESCRIPTION OF THE INVENTION

The invention includes several aspects related to the  
30 display of content to an observer. For example, the invention can enable monitoring of the display of content by a computer system. ~~In particular, the invention can enable monitoring of~~

*insd4*  
*d4*

~~the displayed content in a manner that provides monitoring information from which aspects of the user's observation of the content can be gleaned.~~ The invention can also enable monitoring of the display and - using the aforementioned  
5 capability - the observation at a content display site of content that is provided by a content provider site over a network to the content display site. Additionally, the invention can enable the expeditious provision of updated and/or tailored content over a network from a content provider  
10 site to a content display site so that the content provider's current and appropriately tailored content is always displayed at the content display site.

Herein, "content" refers generally to any sensory images (or data used to produce those sensory images) displayed by a  
15 device with which the invention is used. "Observation" refers to the perception of content by an observer. Typically, the content will be visual or aural images produced by the device; observation of such content thus consists of viewing or listening, as appropriate, to the produced images.

20 Certain aspects of the invention relate to the monitoring of content obtained from, or provision of content over, a network. "Content provider site" refers to a device that is part of the network and that can provide content to another device that is part of the network. "Content display site"  
25 refers to a device that is part of the network and that can receive and display content from another device that is part of the network. It is contemplated that the invention can be particularly useful with a computer network that operates in this way. "Computer network" includes any collection of  
30 interconnected computer systems. "Computer system" refers to a device or collection of devices which depend upon a computational device (e.g., a general or special purpose processor) for at least some aspects of their operation. In

particular, as used herein, a "computer system" can include any type of display device, including a conventional computer display monitor, a television, or one or more audio speakers.

FIGS. 3A, 3B and 3C are simplified diagrams of a network illustrating operation of one aspect of the invention. A content display site 302 (which can be embodied by a conventional client computer) is linked via a network communication line (or lines) 303 to a content provider site 301 (which can be embodied by a conventional server computer). (Typically, the network links multiple content display sites with multiple content provider sites; a single content display site 302 and a single content provider site 301 are shown in FIGS. 3A, 3B and 3C for simplicity. Additionally, it is to be understood that each site on the network can function as both a content display site and a content provider site.) As shown in FIG. 3A, the client computer at the content display site 302 requests content from the server computer at the content provider site 301 over the network communication line 303. As shown in FIG. 3B, the server computer at the content provider site 301 provides content to the client computer at the content display site 302 over the network communication line 303. According to this aspect of the invention, in response to the request for content from the content provider site 301, a set of monitoring instructions (which can be embodied, for example, in a computer program) are also transferred to the content display site 302. The transfer of the monitoring instructions can occur before, with or after the transfer of the content. As explained in more detail below, the monitoring instructions cause the client computer at the content display site 302 to monitor the display of the content to produce monitoring information regarding the manner in which the content is displayed. As shown in FIG. 3C, the monitoring information is transferred from the content display

shea  
ins d5  
10 site 302 to the content provider site 301 over the network  
communication line 303. (The monitoring information could,  
alternatively or additionally, be transferred to another site  
that is part of the network.) <sup>d5</sup> Review of the monitoring  
5 information produced by the monitoring instructions can enable  
conclusions regarding the user's observation of the content to  
~~be deduced, as explained in more detail below.~~ (It should be  
noted that, more generally, monitoring instructions according  
to the invention can be used to monitor the display of content  
10 on a computer system whether or not the computer system is part  
of a network and receives content and monitoring instructions  
over the network.)

00543  
The invention can be used with both public computer  
networks (e.g., the Internet) and private computer networks  
15 (e.g., commercial online services such as America Online™,  
Prodigy™ and CompuServe™, as well as intranets). In  
particular, the invention can be advantageously used with  
computer networks or portions of computer networks over which  
video and/or audio content are transferred from one network  
20 site to another network site for display. Further, the  
invention can advantageously be used with a network in which  
the network sites can be accessed in real time with a browser.  
("Browser" can refer to a computer program which can interpret  
hypertext files and display the content corresponding to those  
25 files, as well as enable transfer from one hypertext file to  
another via a hyperlink within the hypertext file being  
transferred from.) The World Wide Web portion of the Internet  
is a well-known current example of such a network with which  
the invention can be used. Below, some aspects of the  
30 invention are described, for purposes of illustration, as  
implemented in a manner that is compatible with the World Wide  
Web, i.e., in accordance with the hypertext markup language  
(html) and the hypertext transfer protocol (http). However,

none of the aspects of the invention are limited to such implementation.

When the invention is used with a computer network or to monitor display of content by a computer system, aspects of the invention can be implemented as one or more computer programs that can be executed by a computer to achieve the functionality of that aspect. Generally, such computer programs can be implemented using any appropriate computer programming language. However, when an aspect of the invention is used with a computer network that includes computers of many different types (such as the Internet), the computer programming language is preferably one that can be executed by any type of computer (i.e., the computer programming language is platform independent). The Java programming language, developed by Sun Microsystems, Inc. of Mountain View, California, is one such computer programming language. Below, some aspects of the invention are described, for purposes of illustration, as implemented in the Java programming language. Again, however, none of the aspects of the invention are limited to such implementation.

In one embodiment of the invention, the monitoring instructions are transferred to the content display site together with the content. In a particular embodiment, the monitoring instructions are part of a computer program that also includes instructions for displaying the content. Illustratively, such a computer program can be an applet written in the Java programming language. As will be appreciated by those skilled in the use of html, Example 1 below illustrates a set of instructions in accordance with the html syntax that can be used to cause execution of an applet that both displays content and monitors the display.

Example 1

```
<applet code="AdInsert.class" width=230 height=33>  
<param name="image" value="images/southwest.gif">  
<param name="href" value="http://www.swa.com/">  
5 </applet>
```

The instructions shown in Example 1 are executed by a conventional browser implemented on a computer at a content display site when an observer at the content display site makes a request for (e.g., selects a hyperlink) the content represented by the file "southwest.gif." The request is received by an http daemon at the appropriate content provider site. The instructions identify the location ("image") at the content provider site of an applet (a small application program) called "AdInsert" that includes further instructions which, when executed, perform a monitoring method according to the invention, as well as cause the content to be displayed. (The steps that can be implemented in such a monitoring method are discussed further below.) Upon receipt of the request by the http daemon at the content provider site, the AdInsert applet is transferred to the requesting content display site and begins executing. The instructions in Example 1 also establish the size of the area (width and height) in which the content is displayed on a computer display screen, as well as indicate a network site ("href") to which connection can be made by selecting a hyperlink within the content. Thus, illustratively, in accordance with the invention, content from a content provider site that can be accessed by a browser (such as a network site that is part of the World Wide Web) can be transferred to and displayed at a content display site by transferring an applet to the content display site that can be executed by the browser to both display the content and cause aspects of the display of the content to be monitored. (Note that the content being monitored can comprise all of the

content being displayed, or only a part of the content being displayed, e.g., an advertisement present in a Web page.)

In contrast, previously, only content has been transferred to content display sites, using an html syntax as shown in Example 2 below for the content that is displayed by the html syntax shown in Example 1 above.

Example 2

```
<a href="http://www.swa.com/">  
  
10 </a>
```

Thus, previously, it has not been possible to monitor content transferred from a content provider site on the World Wide Web once the content has been transferred to the content display site. As can be appreciated, then, this aspect of the invention provides a powerful tool, not previously available, for obtaining information about how content is displayed on a computer network such as the World Wide Web.

Implementation of a monitoring method as described immediately above means that the operation of the monitoring method is coincident with the display of the content to be monitored. Since the monitoring method does not operate when the content is not being displayed, the monitoring method expends processing capability of the computer system at the content display site only when necessary. At the same time, operation of the monitoring method at all times when the content is displayed is assured.

Further, since the monitoring method can be implemented as part of a broader method according to the invention that also causes the content to be displayed, the problems previously noted with monitoring the display of content that is cached at the content display site are overcome. This is because, unlike the previous use of log files - which require that a request be

made to a content provider site in order that any information be recorded - a monitoring method according to the invention can record monitoring information any time that the content is displayed, without regard to the manner in which the display is requested. In particular, the invention can enable the number of times that a particular set of content is displayed to be precisely counted. This is a huge improvement over previous methods as described above, which not only do not count the number of times that the content is displayed (they count requests), but may not even count the number of requests accurately.

The instructions that implement a monitoring method according to the invention can be used to obtain a large variety of monitoring information. For example, the contents of conventional log files (discussed above) can be ascertained by a monitoring method of the invention. An important aspect of the invention, however, is that monitoring information beyond that available in a conventional log file can also be obtained by a monitoring method according to the invention. Instructions for obtaining several types of such monitoring information are described below. However, it is to be understood that the descriptions below are merely illustrative of the types of monitoring information that can be obtained; the obtaining of other types of monitoring information is also contemplated by the invention.

For example, a monitoring method according to the invention can detect each time that the content is displayed. In fact, in one embodiment of a monitoring method according to the invention, the monitoring method does no more than this. A monitoring method that detects the display of the content can be implemented by an applet as described above. The "monitoring instructions" of such an applet may be no more than an instruction that causes an indication that the applet has

executed to be stored or transferred to an appropriate network site (discussed further below). A monitoring method that can ascertain whether the content was displayed or not can be a very useful monitoring method that provides important basic  
5 information not previously available in an interactive browsing environment for acquiring and viewing content. In particular, undercounting (due to, for example, caching of content at the content display site) and overcounting (due to, for example, submission of artificial requests for content that do not  
10 result in the display of content) of the number of times that the content is displayed are avoided.

A monitoring method according to the invention can also determine the duration of the content display. For example, the duration of the content display can be determined as the  
15 amount of time that the computer program for displaying the content executed, as indicated by time stamps - ascertainable, for example, using a method that exists as part of the Java language - associated with a predefined beginning and end of execution of the program.

20 In one embodiment of the invention, a monitoring method monitors the position of the content on a display screen while the content is being displayed. The position is evaluated to produce monitoring information regarding the display of the content. Such evaluation can be accomplished, for example, by  
25 further obtaining information regarding the position of one or more other images on the display screen, and comparing the position of the content to the position of the one or more other images.

For example, in accordance with the above embodiment, the  
30 monitoring method can determine whether the content is not visible on the display screen, either because the content is occluded by the one or more other images, or because the content has been "scrolled" off of the display screen

(hereafter, these two situations will be referred to together by saying that the content is "hidden"). Further, the monitoring method can determine whether the content is partially hidden, i.e., either partially occluded by the one or  
5 more other images, or partially scrolled off of the display screen. Moreover, the duration of time of each period during which the content is fully or partially hidden can be determined as such periods occur. The duration of time of unobstructed displays of the content can be determined as times  
10 when the view of the content is not either fully or partially hidden. Each of the durations can be reported directly and/or the total duration that the content is fully hidden, partially hidden, fully or partially hidden and/or unobstructed, respectively, can be reported.

15 Information regarding whether or not the displayed content is hidden can be useful for a variety of reasons. For example, such information indicates the amount of time that the displayed content was visible to the observer for observation. Additionally, this information can be used by the content  
20 provider to determine in which regions of a display screen his content is most likely to be unobstructed.

Whether the displayed content is hidden can be determined in any manner that is possible using the tools (e.g., supported programming language, operating system characteristics)  
25 associated with the computer network with which the invention is being used. One way of determining whether the displayed content is hidden using the above-described applet is to periodically declare that the content display (or a portion thereof) is invalid, i.e., the operating system is asked to  
30 redraw the content display, if necessary. If the operating system then makes a request to the applet to redraw the content display, then the content display is not hidden. However, if the operating system does not make a request to the applet to

redraw the content display, then the content display is hidden.

The most complete information regarding whether the content display is hidden can be obtained by invalidating each discrete element of the content display (e.g., pixel) and determining whether the discrete element is hidden, in the manner described above. However, such an approach is computationally expensive and is generally not necessary to obtain useful and sufficiently accurate information regarding whether the content display is hidden. Preferably, then, only a portion of the content display, strategically selected, is evaluated in this manner.

For example, in one embodiment of the invention, each of the corners of the content display are invalidated and monitored for redrawing as described above. If all of the corners are redrawn, then the content display is determined to be unobstructed. If none of the corners are redrawn, then the content display is determined to be fully hidden. If at least one, but not all, of the corners are redrawn, then the display is determined to be partially hidden.

FIGS. 4A, 4B and 4C are simplified views of a display screen 400 including a content display 401 and other images 402, 403 and 404, illustrating an unobstructed, fully hidden, and partially hidden content display, respectively. In FIG. 4A, none of the corners 401a, 401b, 401c or 401d are covered by one of the other images 402, 403 and 404. Thus, after the corners 401a, 401b, 401c and 401d are invalidated, each is redrawn, and the content display 401 is (correctly, in this case) determined to be unobstructed. In FIG. 4B, each of the corners 401a, 401b, 401c and 401d is covered by the image 402. Thus, none of the corners 401a, 401b, 401c and 401d are redrawn after being invalidated, and the content display 401 is (again, correctly) determined to be fully hidden. In FIG. 4C, the corners 401c and 401d are covered by

the image 402, but the corners 401a and 401b are not. Thus, only the corners 401a and 401b are redrawn after invalidation of the corners 401a, 401b, 401c and 401d, and the content display 401 is (once again, correctly) determined to be  
5 partially hidden.

The above approach may not be accurate in all cases. FIG. 4D is a simplified view of the display screen 400 including the content display 400 and other images 402, 403 and 404, illustrating a partially hidden content display 401  
10 that would be determined to be fully hidden according to the method of the invention detailed above. In FIG. 4D, the image 402 covers the corners 401c and 401c, the image 403 covers the corner 401a, and the image 404 covers the corner 401b. Thus, since none of the four  
15 corners 401a, 401b, 401c and 401d is redrawn after being invalidated, the content display 401 is determined to be fully hidden; however, as can be seen in FIG. 4D, this is not the case.

This problem can be alleviated by evaluating other  
20 discrete elements of the content display in addition to the corners. For example, discrete elements at the center of the upper and lower edges, and/or the right and left edges of the content display could be evaluated in addition to the corner pixels. FIG. 4E is a simplified view of the display screen 400  
25 as shown in FIG. 4D, illustrating how another method of the invention can correctly determine the content display 401 to be partially hidden. In FIG. 4E, the upper edge center 401e and the lower edge center 401f of the content display 401 are also evaluated. The lower edge center 401f is covered by the  
30 image 402, while the upper edge center 401e is not. Thus, after invalidation of the corners 401a, 401b, 401c and 401d, and the centers 401e and 401f, the upper edge center 401e is redrawn, and the content display 401 is (correctly) determined

to be partially hidden.

FIG. 4F is a simplified view of the display screen 400, including the content display 401 and other images 402, 403 and 404, illustrating a partially hidden content display 401 that may be determined to be unobstructed according to a method of the invention. In FIG. 4F, none of the corners 401a, 401b, 401c and 401d, or the centers 401e and 401f are covered by the images 402, 403 and 404. Thus, after invalidation of the corners 401a, 401b, 401c and 401d, and the centers 401e and 401f, each is redrawn, and the content display 401 is determined to be unobstructed. However, as can be seen in FIG. 4F, this is not the case, since the image 403 is positioned in the middle of the content display 401.

As illustrated in FIGS 4D through 4F, while the evaluation of additional discrete elements of a content display does not eliminate the possibility of an inaccurate determination regarding whether the content display is hidden, it does reduce the likelihood of such occurrence. Generally, any number and configuration of discrete elements of a content display can be evaluated to reduce the possibility of an incorrect determination regarding whether the content display is hidden, so long as the associated computational cost does not become unacceptably high. Further, the above-described method for determining whether a content display is hidden is only one way in which such determination can be made.

As part of determining whether the content display is hidden, a time stamp is recorded each time there is a change in the "hidden state" of the content display. From these time stamps, the duration of each period of time that the content display is unobstructed, partially hidden and fully hidden can be determined. From the duration of each period, total durations of time that the content display is unobstructed, partially hidden and fully hidden can also be determined.

It may also be possible, by appropriately configuring the discrete elements of the content display that are evaluated, to determine (though, typically, approximately), when the content display is partially hidden, the amount of the content display that is visible. It can be possible, too, when the content display is partially hidden, to give a qualitative description of the portion of the content display that is hidden (or visible), e.g., upper right corner, lower left corner.

When the monitoring method operates on a computer system having an event-driven operating environment, the monitoring method can monitor events as transmitted by the operating system to ascertain information regarding the content display. When the monitoring method is implemented as an applet that also displays the content, such monitoring can occur naturally, since only events concerning the content display are transmitted to the monitoring method. For example, the applet can use a pre-existing Java method (e.g., the method named `HandleEvent` in a current version of Java) to monitor events as transmitted by the operating system. Such event monitoring can be used to, for example, determine the number of times that an on-screen pointer (e.g., a mouse arrow or a cursor) entered an area defined by the content display. (The defined area can be related to the content display in any manner and can be, for example, the area in which the content is displayed, or an area somewhat smaller or larger than the area of the content display.) The operating system of the computer system displaying the content display typically monitors the position of the on-screen pointer and can identify in which region on the display screen the pointer is located. Thus, an applet configured to display content, as described above, can discern whether the pointer is located within the content display by monitoring an event that indicates that the pointer has entered the area defined by the content display. The monitoring method

of the invention can use this information provided by the operating system to count the number of times that the on-screen pointer enters the area defined by the content display. The monitoring method can also determine when the on-screen pointer leaves the defined area after each entry, by monitoring another event that indicates that the pointer has exited the area defined by the content display. The time stamps associated with the entry into and exit from the defined area can be used to calculate the duration of time that the pointer was in the defined area for each entry into the defined area, as well as the total duration of time that the pointer was within the defined area. The monitoring method can also determine when the on-screen pointer is moving within the defined area, again by monitoring an event that indicates such pointer movement. The above-described information regarding the on-screen pointer position and movement relative to the content display may be useful in determining how attentive the observer was to the content, since an observer frequently watches the position of the on-screen pointer when viewing the display screen.

In another embodiment of the invention, the monitoring method monitors the change in time of a characteristic of the content display. The change in time of this characteristic is evaluated to produce monitoring information. The evaluation can be accomplished, for example, by further monitoring the change in time of a characteristic of the computer system used to display the content, and comparing the change in time of the characteristic of the content display to the change in time of the characteristic of the computer system. Either of the two examples given immediately above (hiding of the content display and entry of a pointer into a defined area) are also examples of a monitoring method in accordance with this embodiment of the invention.

A monitoring method according to the invention can obtain a variety of other information, as well. For example, the monitoring method can obtain a time stamp (date and time of day) that indicates when the display of the content began.

5 When the monitoring method is implemented by an applet written in Java, the time stamp can be obtained using a method that exists as part of the Java language.

Identifying information regarding the computer on which the content is displayed can also be obtained. The Internet

10 Protocol (IP) address from which the request for the content was made, as well as an identification of the machine to which the content was transferred can be obtained. (There may not be a one-to-one correspondence between these two if, for example, the latter is a client computer of a system for which the

15 former is a server computer.) Again, both the IP address and machine name can be obtained using a pre-existing Java method.

A monitoring method according to the invention can also determine if the user of the computer at the content display site selected (e.g., clicked with a mouse or pressed an

20 appropriate keyboard key) a hyperlink within the area of the content display to end display of the current content display. Similar to the monitoring of the pointer location described above, an applet that implements a monitoring method of the invention can include a standard Java method (e.g.,

25 HandleEvent) that accepts events transmitted by the operating system. One of the events is the selection of a hyperlink. When such an event is reported, the monitoring method can so note.

As previously indicated, the above-described examples of

30 monitoring information are merely illustrative of the types of monitoring information that can be obtained by a monitoring method according to the invention. Generally, a monitoring method according to the invention can make use of any method

available in the computing environment, e.g., an operating system method, or a method that is part of a software framework, or that can be written in a computer programming language that can be used in the computing environment. For example, when the monitoring method is implemented by an applet written in Java, any existing Java method can be used to obtain information relevant to the display of the content to be monitored, either by using the method to change the state of the computer (e.g., the state of the display) on which the content is being displayed and monitoring the response of the computer (e.g., the method for monitoring whether content display is hidden, discussed above) or by retrieving information about the state of the computer (e.g., the method for monitoring entry of the pointer into the content display, discussed above). In particular, the monitoring of events as discussed above can be useful in discerning information about the content display.

A monitoring method according to the invention can also be used to ascertain information about an audio display. For example, if the content being monitored includes audio content that can only be displayed by selecting an appropriate user interface mechanism (e.g., a graphical pushbutton), a monitoring method according to the invention can determine whether that "event" is transmitted to the window represented by the content, indicating that the audio display was at least begun. Using a method as described above for determining the duration of a content display, together with knowledge of the when the audio display was begun (using a time stamp as described above), the duration of the audio display can also be determined. It may also be possible to determine the volume at which the audio content is displayed, by appropriately monitoring the methods used to operate the audio display devices. These examples are merely illustrative. As can be

appreciated, using any other method available in the computing environment, other information regarding an audio display can be determined.

A monitoring method according to the invention can also be  
5 used to explicitly (i.e., by presenting questions to observers  
that they can answer) acquire demographic information regarding  
the observers of the content being monitored. This could be  
implemented, for example, by including the instructions for  
asking such questions, the content of the questions and the  
10 instructions for storing the obtained demographic information  
in a computer program used to implement the monitoring method.  
Or, such instructions and question content could be stored in  
a separate file that is called and executed by the computer  
program that implements the monitoring method. Or,  
15 instructions for presenting the questions and storing the  
answers could be included as part of the computer program for  
implementing the monitoring method, and the content of the  
questions could be contained in a separate file that is  
accessed by the computer program. These latter two  
20 possibilities can be particularly advantageous, since they  
allow multiple sets of demographic questions to be presented to  
observers by the monitoring method, thus enabling the  
demographic questions to be tailored to the content being  
displayed or to the characteristics of the observer likely to  
25 view the content.

As described above, in accordance with the invention,  
monitoring information regarding the display of content can be  
obtained, then later reviewed and analyzed to enable  
conclusions to be drawn about how the content was displayed  
30 and, possibly, to enable deductions to be made about how the  
content was observed. In addition, monitoring information can  
be used to affect the display of a set of content. One way in  
which this can occur is for a set of content, or the manner in

which the set of content is displayed, to be modified based upon review and analysis of monitoring information obtained from previous displays of the set of content (e.g., monitoring information regarding whether or not the content was hidden, or  
5 the frequency of display of the content at different times during the day or week, that may be used to determine the best location on a display screen or the best times, respectively, to display the content).

Another way in which monitoring information can be used to  
10 affect the display of a set of content is to use certain monitoring information obtained just before or during the display to cause the set of content to be displayed in a particular manner. For example, as discussed above, the IP address from which a request for a set of content emanated can  
15 be ascertained when the request is first received. It may be possible to associate characteristics of an observer or observers with an IP address from which a request for content has been received (because, for example, demographic information has previously been obtained as described above  
20 when a set of content was previously transferred to that IP address). Based upon the known characteristics associated with the IP address, an appropriate one of multiple versions of the requested set of content can be transferred for display, e.g., if it is known that the IP address corresponds to a content  
25 display site that is used by observers that speak a particular language, then text displayed aurally or visually can be displayed in that language. As another example, the duration of time that a set of content has been displayed can be determined, as discussed above, and the portion of the set of  
30 content that is being displayed changed as a function of that duration, e.g., the display of a set of content can begin with a particular video display and change to another video display after passage of a specified duration of time. As still

another example, the portion of a set of content that is displayed can be varied based upon performance characteristics of the network over which the content is transferred. For example, the amount of time required to transfer data from the content provider site to the content display site can be monitored (by, for example, obtaining information from the log file regarding the size of transferred files and the amount of time required to transfer those files, as discussed above). The display of the content can then be controlled so that a moving video is displayed if the data transfer rate is above a predetermined magnitude and a still video is displayed if the data transfer rate is below the predetermined magnitude, the predetermined magnitude being chosen so that data rates below that magnitude are insufficiently fast to produce moving video of acceptable quality. The above examples are merely illustrative; other ways of using monitoring information to affect the display of a set of content are contemplated by the invention.

As described above, a monitoring method according to the invention can obtain monitoring information regarding the display of content. Of particular interest is the basic question of whether the content was displayed at all. As described above, a monitoring method according to the invention can make this determination. Some observers, however, have developed techniques for suppressing the display of particular content (e.g., advertisements). A monitoring method according to the invention can also increase the likelihood that particular content is displayed by conditioning certain other operation of the computer system that displays the content on the display of that particular content. For example, the content to be monitored can be presented as part of other content. Such presentation is common on, for example, the World Wide Web, where, for example, advertising content is

frequently included as part of other content. A monitoring method according to the invention can condition the display of the other content on the display of the to-be-monitored content, e.g., the full content of a Web page cannot be viewed unless an advertisement included on the Web page is viewed. Moreover, the presentation of the other content can be conditioned on the display of the to-be-monitored content for a specified period of time. This can be particularly valuable when the to-be-monitored content does not appear automatically as part of the other content, but, rather, is only displayed in response to selection of an appropriate user interface mechanism (e.g., a graphical pushbutton) that is part of the other content.

In accordance with the above-described forced presentation of specified content, the detection of content suppression can be accomplished in any suitable manner. For example, it may be possible to detect the suppression technique being used. <sup>or,</sup> the display of the content can be ordered so that the content that must be displayed ("required content") is displayed first; if the monitoring method detects that the content display site is displaying the other content without first having displayed the required, then suppression of the required content has been detected. Upon detection of suppression of the required content, display of the other content is prevented and, if desired, a message indicating that fact can be displayed.

As indicated above, after monitoring information is obtained by a monitoring method according to the invention, the monitoring information is transferred from the content display site 302 to a remote site. The remote site can be the content provider site 301 or another site that is part of the network. When a monitoring method according to the invention is implemented by a Java applet, the remote site is the content provider site 301, since, currently, such applets can only

communicate information to the network site from which they were transferred. However, in the future, such constraint may not exist; in that event, the remote site need not necessarily be the content provider site 301 even when a Java applet is used to implement a monitoring method according to the invention.

Generally, the monitoring information can be transferred to the remote site at any time. It may be desirable, for example, to cause the monitoring information to be transferred to the remote site immediately after the monitoring information is obtained, so that the monitoring information is accessible as quickly as possible. It may, alternatively, be desirable to store the monitoring information at the content display site, then transfer the monitoring information at a time when communication over the network communication line 303 is fastest and/or least expensive, e.g., at night.

The monitoring information can be communicated to a communication port that is different than the port from which the content and the monitoring instructions were transmitted to the content display site 302. In that event, a special daemon that monitors such communication port for receipt of monitoring information is installed on the server computer at the content provider site. The daemon can be implemented as a conventional server daemon for monitoring data received by a server computer on a designated communication port. Communication of the monitoring information to a specially designated port can be useful to enable the monitoring data to sent in any desired format in accordance with any desired protocol. For example, the monitoring data can be encrypted, as described below.

When the invention is implemented with the World Wide Web, it is also possible to transmit the monitoring data over the network using the communication channel monitored by the http daemon, i.e., by transmitting a request to the http daemon.

Such transmission may be desirable for several reasons. For example, transmission of monitoring data to the http daemon eliminates the need to create, and supply to operators of the remote site to which the monitoring data is to be transferred  
5 (e.g., a Web page operator or, as described below, an application manager site operator), special software for receiving the monitoring data. Additionally, transmission of monitoring data by transmitting a request to the http daemon may be the only way to transfer the monitoring data to the  
10 remote site. This can be true, for example, when one or more client computers are served by a "proxy server" which mediates communication between the client computers and other sites on the network. The proxy server may not allow communication over a channel specially designated for transmitting monitoring  
15 data, but allow communication to the http daemon.

Transmission of monitoring data by making a request to the http daemon can be accomplished in a variety of ways. For example, an http request can be submitted for a file having a "name" that denotes the monitoring data in some way.  
20 Notwithstanding the spurious nature of the file request, the request is recorded in the http log file, from which the "name" can be retrieved to enable extraction of the monitoring data. Or, a request for execution of a CGI script can be transmitted, with the parameter of the CGI script request that specifies  
25 input to the script being specified to denote the monitoring data in some way. A computer program resident on the computer system at the remote site can then implement a method that extracts the value of the input from the CGI script, and the monitoring data can be extracted from the value of the input.  
30 Other methods of using CGI scripts or http requests to transmit monitoring data to an http daemon are possible.

For security, it may be desirable to encrypt the monitoring data before it is transferred from the content

display site 302 to a remote site. Any suitable encryption method can be used. For example, a public key encryption method, such as the well known RSA algorithm, can be used to encrypt the monitoring data. In general, the monitoring data  
5 (or other data transferred over a network in accordance with the invention) can be encrypted before any transmission of the data over a network (other examples of such data transmission are described below as part of the systems illustrated in FIGS. 5A, 5B and 5C, and FIGS. 6A, 6B, 6C and 6D).

10 Once communicated to the remote site, the monitoring information can be stored in any appropriate database, as known to those skilled in the art of constructing and managing databases. The monitoring information can be presented for observation through a suitable user interface, such as a  
15 graphical user interface (GUI), in any desired format, e.g., graphs, bar charts, pie charts. The monitoring information stored in the database can also be subjected to further analysis if desired. For example, the total time that a content display is available to be viewed can be broken down  
20 into percentages of time that the content display was unobstructed, partially hidden and fully hidden. Or, the percentage of observers of a set of content that select a particular hyperlink while observing the content can be identified.

25 The monitoring information may be of interest not only to the content provider that provides the content for display, but to third parties as well. For example, if the content provided by the content provider includes an advertisement, the advertiser may be interested in the monitoring information  
30 regarding display of the content. The third party and the content provider may have conflicting interests in the substance of the monitoring information. For example, if the third party is paying the content provider to include the third

party's content with the content provider's content, and the payment is based upon the amount of exposure of the third party's content to observers, the content provider has an interest in the monitoring information showing a large amount of exposure of the content, while the third party has an interest in the monitoring information showing a small amount of exposure. (Both parties, of course, can be simultaneously motivated by other interests, as well: for example, the third party may simply want the monitoring information to reflect accurately the amount of exposure of the content, so that they can use that information in assessing the effects of providing their content through the content provider.) If the monitoring information is transferred from the content display site to the content provider site, and unrestricted access to the monitoring information allowed at the content provider site, there may be no foolproof way to prevent the content provider from tampering with the substance of the monitoring information. This problem is particularly acute when a monitoring method according to the invention is embodied in a manner (e.g., by a Java applet), as discussed above, that necessitates that the monitoring information be transferred back to the content provider site.

FIGS. 5A, 5B and 5C are simplified diagrams of a network illustrating operation of another embodiment of the invention. This embodiment of the invention provides a system configuration that can overcome the problem of possible tampering with the substance of the monitoring information by the content provider. As in the system illustrated in FIGS. 3A, 3B and 3C, a content display site 302 is linked over a network to a content provider site 301. The network also includes an application manager site 501. The content display site 302 and content provider site 301 can communicate with each other via the network communication line 303, as described

above, to enable transfer of content and monitoring instructions from the content provider site 301 to the content display site 302. Alternatively, the content and monitoring instructions can be transferred to the content display site 302 from the application manager site 501 in response to a request received from the content provider site 301 upon receipt of the request from the content display site 302. This latter implementation is illustrated in FIGS. 5A, 5B and 5C. In such an implementation, the content provider site 301 needn't have either a computer program for implementing the monitoring method or a program for receiving monitoring data installed at the content provider site 301, thus simplifying use of the invention for the content provider. Rather, the content provider need only have an open account (as discussed below) at the application manager site 501.

In this embodiment of the invention, the monitoring information obtained at the content display site 302 is transferred to the application manager site 501, either directly from the content display site 302 or indirectly via the content provider site 301. If the latter, then the monitoring information can be received by the content provider site 301 and transferred to the application manager site 501 in a way that prevents access to the monitoring information at the content provider site 301. For example, the monitoring information could be encrypted at the content display site 302 before transfer to the content provider site 301, the decryption method being available only at the application manager site 501. Or, the monitoring information could be immediately transferred to the application manager site 501 after being received at the content provider site 301. Once received at the application manager site 501, access to the monitoring information can be administered by the (neutral) application manager so that the monitoring information can not

be modified by any of the parties having an interest in the information, thus ensuring the integrity of the monitoring information.

In a typical implementation, multiple sets of content will be provided from multiple content provider sites, and each set of content will be displayed by multiple content display sites. A set of monitoring information will be recorded for each display of each of the multiple sets of content and transferred to the application manager site for storage in a database that is implemented on a computer at the application manager site. Each set of monitoring information must be identified as corresponding to the set of content for which the monitoring information was obtained, so that monitoring information can be appropriately stored in a database to enable later retrieval of the monitoring information for that set of content. When a monitoring method according to the invention is implemented for use with the World Wide Web, this can be accomplished by appropriate specification of a parameter included in a computer program written in html used to implement a monitoring method, as discussed above. Example 3 below illustrates how Example 1 discussed above can be modified to make such specification (the "Account" parameter).

### Example 3

```
<applet code="//AppMgr.com/AdInsert.class" width=230 height=33>  
25 <param name="image" value="images/southwest.gif">  
  <param name="href" value="http://www.swa.com/">  
  <param name="Account" value="9004560093">  
</applet>
```

The database residing on the computer at the application manager site can also be used, for example, to store account information about the content provider site from which the content display is provided.

~~In the embodiment of the invention illustrated in FIG. 5,~~

5 ~~a user interface (e.g., GUI) can be provided on the content~~  
provider site computer to enable the owner (or representative)  
of the content provider site to ~~access monitoring information~~  
stored at the application manager site regarding content  
10 ~~displays provided by the content provider site.~~ Such an  
interface can also be configured to enable the content provider  
to create a new account on the application manager computer,  
authorize payments for use of the monitoring system of the  
invention, and request particular analysis or presentation of  
10 obtained monitoring information. Other functions can also be  
provided in such an interface, as desirable.

It is also possible that there be multiple application  
manager sites. Typically, monitoring information for each  
content display will be designated for storage on a particular  
15 one of the application manager sites. Such designation can be  
included as a parameter specification in a computer program  
used to implement the monitoring information as discussed  
above.

As discussed above, the content provided by a content  
20 provider can be tailored according to any specified criteria.  
Further, the content provider may periodically update the  
content. Additionally, third parties may want to provide their  
content with that of a content provider. These third parties  
may also have multiple sets of specially tailored content that  
25 are updated periodically. The management of such multiple sets  
of content by a content provider at the content provider site  
can become undesirably complex and may overtax the available  
bandwidth for transmission of data to and from the content  
provider site.

30 FIGS. 6A, 6B, 6C and 6D are simplified diagrams of a  
network illustrating operation of still another embodiment of  
the invention. This embodiment of the invention provides a  
system configuration that can enable updated and/or tailored

secondary content provided by a secondary content provider to  
be transferred to a primary content display site for use with  
primary content supplied by a primary content provider without  
the problems identified above with existing such systems, as  
discussed in more detail below. ~~This embodiment of the  
invention also can enable all of the functionality described  
above for the system illustrated in FIGS 5A, 5B and 5C.~~ In the  
embodiment of the invention shown in FIGS. 6A, 6B, 6C and 6D,  
a content display site 302, a primary content provider site 602  
and an application manager site 501 are linked to each other  
over a network and can communicate with each other as described  
above. The network also includes a secondary content provider  
site 601. As shown in FIG. 6A, in this embodiment of the  
invention, secondary content can be provided from the secondary  
content provider site 601 to the application manager site 501  
and stored thereat. As shown in FIG. 6B, whenever secondary  
content is provided to the application manager site 501, the  
application manager causes the content to both be stored at the  
application manager site 501 and transferred to all content  
provider sites, e.g., content provider site 602, that provide  
that secondary content with their primary content. When a  
request for the primary content that includes such secondary  
content is received by the primary content provider site 602  
from the content display site 302 (FIG. 6C), the primary  
content provider site 602 is able to immediately (i.e., without  
necessity to retrieve the content from another network site or  
request that the content be provided to the content display  
site from another network site) provide both the primary and  
secondary content to the content display site 302, as shown in  
FIG. 6D.

As can be appreciated, the management of both primary and  
secondary content can become quite burdensome when many sets,  
and/or many versions of sets, of secondary content and/or

primary content are being provided from the primary content provider site. Management of continual updates to these sets of content data exacerbates this burden. By storing secondary content (and, if desired, primary content) at the application manager site 501, the system of FIGS. 6A, 6B, 6C and 6D relieves the primary content provider of the burden of managing such content. However, because the application manager causes the content to be stored at the content provider site 602, the content can be provided to the content display site 302 from the content provider site 602, rather than the application manager site 501, thus leaving control of bandwidth management with the primary content provider so that the primary content provider can ensure that a system that adequately addresses the bandwidth requirements of the content provided from the primary content provider site 602 is in place. This is an important consideration for the primary content provider since requesters of content from the primary content provider will hold the primary content provider responsible for the performance characteristics (e.g., speed) associated with the provision of that content. The system of FIGS. 6A, 6B, 6C and 6D, then, relieves the primary content provider of the need to manage the storage of content, while reserving control over the provision of that content to the content provider.

In this embodiment of the invention a user interface (e.g., GUI) can be provided at both the primary content provider site 602 and the application manager site 501. <sup>the</sup> ~~The primary content provider user interface can provide the same functionality as described above with respect to FIG 5.~~

Additionally, the primary content provider user interface can enable the content provider to select available secondary content for possible inclusion with that content provider's primary content. Such selection can also include specification of terms upon which the primary content provider wishes to

include the secondary content. Selection of secondary content does not automatically cause the secondary content to be included with the primary content provider's content, but, rather, causes a request for such inclusion to be made (e.g., via the secondary content provider user interface, described below) to the secondary content provider. Upon acceptance by the secondary content provider, the secondary content can be included with the primary content. The secondary content provider user interface can enable the secondary content provider to select a primary content provider site with which to include the secondary content provider's content. Again, such selection can be made together with specification of the terms of such inclusion; the selection causes a request for inclusion to be made (e.g., via the primary content provider user interface) to the primary content provider. ~~The secondary content user interface can also provide functionality similar to that described above with respect to FIG. 5.~~ As will be readily appreciated by those skilled in the art, other functions can also be provided in the primary content provider and secondary content provider user interfaces, as desirable. ~~The embodiment of the invention shown in FIG. 6 facilitates interaction between the primary content provider site 602 and the secondary content provider site 601 to enable a secondary content provider to easily and flexibly provide content to a primary content provider in a manner that enables both the primary and secondary content providers to exercise control over the provision of content.~~

As described above, monitoring instructions and content can be embodied by an applet that executes at the content display site. In the system of FIGS. 5A, 5B and 5C or the system of FIGS. 6A, 6B, 6C and 6D, the use of such an applet can advantageously dovetail with an implementation of those systems in which the applet is transferred to the content

display site from the application manager site. This is because an applet must return the monitoring information to the network site from which the applet was transferred: thus, the monitoring information is transferred directly to the neutral  
5 application manager site. Too, when the monitoring instructions and content are transferred from the application manager site to the content display site, the use of monitoring information to tailor the content provided to the content display site, as discussed above, can also be easily  
10 implemented.

Various embodiments of the invention have been described. The descriptions are intended to be illustrative, not limitative. Thus, it will be apparent to one skilled in the art that certain modifications may be made to the invention as  
15 described above without departing from the scope of the claims set out below. For example, though the invention has been described above as it particularly applies to monitoring the display of content disseminated over the World Wide Web, the invention can generally be used to monitor the display of  
20 content disseminated over any computer network. Additionally, though an implementation of the invention has been described in which aspects of the Java programming language are used, it is to be understood that invention is not limited to such implementation; other programming languages could be used  
25 having other features and characteristics (e.g., the language need not be an object-oriented language as is Java).